

Interpretable clustering using unsupervised binary trees Ricardo Fraiman, Badih Ghattas, Marcela Svarc

▶ To cite this version:

Ricardo Fraiman, Badih Ghattas, Marcela Svarc. Interpretable clustering using unsupervised binary trees. 2024. hal-01297499

HAL Id: hal-01297499 https://hal.science/hal-01297499

Preprint submitted on 11 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Interpretable Clustering using Unsupervised Binary Trees

Ricardo Fraiman^{*}, Badih Ghattas^{**} and Marcela Svarc^{***1}

*Universidad de San Andrés, Argentina and Universidad de la República, Uruguay.
** Université de la Méditerrannée Facult des Sciences de Luminy, Dpartement de Mathmatiques 163 Avenue de Luminy, 13288 Marseille cedex 09, France.
*** Universidad de San Andrés and CONICET, Argentina.

Abstract

We herein introduce a new method of interpretable clustering that uses unsupervised binary trees. It is a three-stage procedure, the first stage of which entails a series of recursive binary splits to reduce the heterogeneity of the data within the new subsamples. During the second stage (pruning), consideration is given to whether adjacent nodes can be aggregated. Finally, during the third stage (joining), similar clusters are joined together, even if they do not share the same parent originally. Consistency results are obtained, and the procedure is used on simulated and real data sets.

<u>Keywords</u>: Unsupervised Classification, CART, Pattern Recognition. <u>Running Title</u>: Clustering a la CART.

1 Introduction

Clustering is a means of unsupervised classification, is a common technique for the statistical analysis of data, and has applications in many fields, including medicine, marketing and economics, among other areas. The term "cluster analysis" (first used by Tryon, [24]) includes the use of any of a number of different algorithms and methods for grouping similar data into a number of different categories. The grouping is achieved in such a way that "the degree of association between data is at a maximum if the data belong to the same group and at a minimum otherwise".

Cluster analysis or clustering involves the assignment of a set of observations from \mathbb{R}^p into subsets (called clusters), such that observations in the same cluster are similar in "some sense". The definitions are quite vague, in that there is no clear objective function of the population that can be used to measure the performance of a clustering procedure. Implicit in each clustering algorithm is an objective function that varies from one method to another. It is important to note that although most clustering procedures require the number of clusters to be known beforehand, generally in practice it is not.

In contrast, in supervised classification the number of groups is known and we additionally have both a learning sample and a universal objective function, i.e. to minimise the number of misclassifications, or in terms of the population, to minimise the Bayes error.

¹Corresponding author: Marcela Svarc, Departamento de Matemáticas y Ciencias, Universidad de San Andrés, Vito Dumas 284, Victoria (1644), Buenos Aires, Argentina. Email: msvarc@udesa.edu.ar

Despite these differences, there are a number of similarities between supervised and unsupervised classification. Specifically, there are many algorithms that share the same spirit in both cases.

Algorithms that use supervised and unsupervised classification [13] can either be partitional or hierarchical. Partitional algorithms determine all the groups at once. The most widely used and well-studied partitioning procedure for cluster analysis is that of the k-means algorithm. Hierarchical algorithms successively identify groups that split from or join groups that were established previously. These algorithms can either be agglomerative ("bottom-up") or divisive ("top-down"). Agglomerative algorithms begin with each element as a separate group and merge them into successively larger groups. In contrast, divisive algorithms begin with the whole set and proceed to split it into successively smaller groups. Hierarchical algorithms create a hierarchy of partitions that may be represented in a tree structure. The best known hierarchical algorithm for supervised classification is CART [4].

CART has a further property of interest. The partition tree is built using a few binary conditions obtained from the original coordinates of the data. In most cases, the interpretation of the results may be summarised in a tree that has a very simple structure. The usefulness of such a scheme of classification is valuable not only for the rapid classification of new observations, but it can also often yield a much simpler "model" for explaining why the observations are classified in a particular group, a property that is remarkably important in many applications. Moreover, it is important to stress that the algorithm assumes no kind of parametric model for the underlying distribution.

Recently, several new methods have been proposed for clustering analysis, see for instance Garcia Escudero et al. [10] for a review with focus on robust clustering procedure. Other recent proposals have been made by Peña and Prieto ([18],[19]), Fraley and Raftery [7], Oh and Raftery [16], Walther [25], among others. But just a few different methods using decision trees to obtain clusters have previously been proposed. Liu et al. [14] use decision trees to partition the data space into clusters and empty (sparse) regions at different levels of detail. Their method uses the idea of adding an artificial sample of size N that is uniformly distributed over the space. With these N points added to the original data set, the problem then becomes one of obtaining a partition of the space into dense and sparse regions. Liu et al. [14] treat this problem as a classification problem that uses a new "purity" function that is adapted to the problem and is based on the relative densities of the regions concerned.

Chavent et al. [5] obtained a binary clustering tree that applies to a particular variable and its binary transformation. They presented two alternative procedures. In the first, the splitting variables are recursively selected using correspondence analysis, and the resulting factorial scores lead to the binary transformation. In the second, the candidate variables and their variable transformations are simultaneously selected by a criterion of optimisation in which the resulting partitions are evaluated. Basak et al. [2] proposed four different measures for selecting the most appropriate characteristics for splitting the data at every node, and two algorithms for partitioning the data at every decision node. For the specific case of categorical data, Andreopoulus et al. [1] introduced HIERDENC, which is an algorithm that searches the dense subspaces on the "cube" distribution of the values presented in the data at hand. Our aim herein is to propose a simple clustering procedure that has the same appealing properties as CART. We introduce the hierarchical top-down method of CUBT (Clustering using unsupervised binary trees), in which the clustering tree is based on binary rules on the original variables, and this will help us to understand the structure of the clustering.

There are three stages in our procedure. In the first, we grow a maximal tree by applying a recursive partitioning algorithm. In the second, we prune the tree using a criterion of minimal dissimilarity. In the final stage, we aggregate the leaves of the tree that do not necessarily share the same direct ascendant.

We do not claim that the new method we introduce is always more efficient nor better than others. For each particular problem, depending on the cluster structure some methods behave better than others, and quite often the difference in efficiency is just a small (but important) improvement. This is also the case in supervised classification. CART is far from being the best universally more efficient algorithm. However it has a quite good behavior for a large class of classification problems. We will show along the paper, that the main advantages of our clustering method (mostly shared with CART for classification problems) are the following:

- a) **Flexibility.** The method is able to perform good clustering for a large family of cluster structure. As long as the true clusters can be separated by some partition built as the intersection of a arbitrarily large finite number of half-spaces, whose boundaries are orthogonal to the original coordinates system the method will work properly.
- b) **Interpretability.** The final partition is explained in terms of binary partitions on the original variables. This property is fundamental for many applications.
- c) Efficiency. We get a good performance in terms of correct clustering allocation for a large family of clusters structures.
- d) **Population version.** We provide a population version of the final partition, regarded as a random vector $\mathbf{X} \in \mathbb{R}^p$ with unknown distribution P. We then show that the algorithm (the empirical version) converges a.s. to the population final partition. This kind of property is essential in statistics in order to understand well when and why a method will be adequate, by looking at the population version. This is briefly discussed on Section 2.3.

The remainder of the paper is organised as follows. In Section 2, we introduce some notation and we describe the empirical and population versions of our method. The latter describes the method in terms of the population, regarded as a random vector $\mathbf{X} \in \mathbb{R}^p$ with an unknown distribution P. The consistency of our method is described in Section 3. In Section 4, we present the results of a series of simulations, in which we used our method to consider several different models and compare the results with those produced by the k-means algorithm, MCLUST and DBSCAN. Using a synthetic data set, we also compared the tree structures produced by CART (using the training sample, with the labels) and CUBT, considering the same sample in each case without the labels. A set of real data is analysed in Section 5, and our concluding remarks are made in Section 6. All proofs are given in the Appendix.

2 Clustering a la CART

We begin by establishing some notation. Let $\mathbf{X} \in \mathbb{R}^p$ be a random p-dimensional real vector with coordinates X(j), j = 1, ..., p, such that $\mathbb{E}(\|\mathbf{X}\|^2) < \infty$. The data consist in n random independent and identically distributed realisations of $\mathbf{X}, \Xi = {\mathbf{X}_1, ..., \mathbf{X}_n}$. For the population version the space is \mathbb{R}^p , while for the empirical version the space is Ξ . We denote the nodes of the tree by t. Each node t determines a subset of \mathbb{R}^p , $t \subset \mathbb{R}^p$. We assign the whole space to the root node.

Even though our procedure uses the same general approach as CART in many respects, two main differences should be stressed. First, because we use unsupervised classification, only information about the observations without labels is available. Thus the splitting criterion cannot make use of the labels, as it can in CART. The second essential difference is that instead of having one final pruning stage, in our algorithm we subdivide this stage into two, in that we first prune the tree and then use a final joining process. In the first of these two procedures, we evaluate the merging of adjacent nodes, and in the second the aim is to aggregate similar clusters that do not share the same direct ascendant in the tree.

2.1 Forward step: maximal tree construction

Because we are using a top-down procedure, we begin by assigning the whole space to the root node. Let t be a node and $\hat{t} = \Xi \cap t$, the set of observations obtained from the sample concerned. At each stage, a terminal node is considered to be split into two subnodes, namely the left and right child, t_l , t_r , if it fulfills a certain condition. At the beginning there is only one node, i.e. the root, which contains the whole space. The splitting rule has the form x(j) < a, where x(j) is a variable and a is a threshold level. Thus, $t_l = \{x \in \mathbb{R}^p : x(j) \le a\}$ and $t_r = \{x \in \mathbb{R}^p : x(j) > a\}$.

Let X_t be the restriction of X to the node t, i.e. $X_t = \mathbf{X} | \{ \mathbf{X} \in t \}$, and α_t the probability of being in $t, \alpha_t = P(\mathbf{X} \in t)$. R(t) is then a heterogeneity measure of t and is defined by,

$$R(t) = \alpha_t \ trace(Cov(X_t)),\tag{1}$$

where, $cov(X_t)$ is the covariance matrix of X_t . Thus, R(t) is an approximate measure of the "mass concentration" of the random vector X at set t, weighted by the mass of set t. In the empirical algorithm α_t and $Cov(X_t)$ are replaced by their empirical versions (estimates), and R(t) is called the *deviance*. We then denote n_t to be the cardinal of the set \hat{t} , $n_t = \sum_{i=1}^n \mathcal{I} \{ \mathbf{X}_i \in t \}$, (where \mathcal{I}_A stands for the indicator function of set A), and hence the estimated probability is $\hat{\alpha}_t = \frac{n_t}{n}$, the estimate of $E(\|X_t - \mu_t\|^2)$ is

$$\frac{\sum_{\{\mathbf{X}_i \in t\}} \left\| \mathbf{X}_i - \overline{X}_t \right\|^2}{n_t},$$

where \overline{X}_t is the empirical mean of the observations on t and the estimate of the deviance is,

$$\hat{R}(t) = \frac{\sum_{\{\mathbf{X}_i \in t\}} \left\| \mathbf{X}_i - \overline{X}_t \right\|^2}{n}.$$
(2)

The best split for t is defined as the couple $(j, a) \in \{1, \ldots, p\} \times \mathbb{R}$, (the first element indicates the variable where the partition is defined and the second is the threshold level) that maximises,

$$\Delta(t, j, a) = R(t) - R(t_l) - R(t_r).$$
(3)

It is easy to verify that $\Delta(t, j, a) \ge 0$ for all t, j, a, and this property is also verified by all the splitting criteria proposed in CART.

Remark 2.1. (Uniqueness) As in CART the maximum in (3) may not by unique. Typically, the maximum is attained only for one variable, j, but there may be many values a for which it is attained, for instance a union of intervals. In those cases we will choose the smallest value a for which the maximum is attained. More precisely, for fixed j, a is defined as

 $\inf \{ \operatorname{argmax}_{a \in \mathbb{R}} \Delta(t, j, a) \}.$

If there are several variables that attain the maximum then we choose the variable j with smallest index, and then if needed we choose the smallest value a where the maximum is attained.

We fix two parameters τ , mindev $\in (0, 1)$.

We begin with the whole space being assigned to the root node, and each node is then split recursively until one of the following stopping rules is satisfied: (i) $\alpha_t < \tau$; (ii) The reduction in deviance is less than $mindev \times R(\mathcal{S})$, where \mathcal{S} is the whole space. For the empirical version of the algorithm we replace α_t by $\hat{\alpha}_t$ and $\Delta(t, j, a)$ by $\hat{\Delta}(t, j, a) = \hat{R}(t) - \hat{R}(t_l) - \hat{R}(t_r)$, and denote by $minsize = [\tau n]$ (the minimum size of a node). minsizeand mindev are tuning parameters that must be supplied by the user. The uniqueness problem also appears on the empirical version, and it can be treated as in the population algorithm.

When the algorithm stops, a label is assigned to each leaf (terminal node), and we then call the actual tree the maximal tree. At this point, we have obtained a partition of the space and in consequence a partition of the data set, in which each leaf is associated with a cluster. Ideally, this tree has at least the same number of clusters as the population, although in practice it may have too many clusters, and then an agglomerative stage must be applied as in CART. It is important to note if the number of clusters k is known, the number of leaves should be greater or equal to k. Small values of *mindev* ensure a tree that has many leaves. Moreover, if the tree has the same number of leaves as the number of clusters, it is not then necessary to run the subsequent stages of the algorithm.

2.2 Backward step: pruning and joining

In the next step, we successively use two algorithms to give the final conformation of the grouping. The first prunes the tree and the second merges non-adjacent leaves (we herein refer to this as "joining"). We now introduce a pruning criterion that we refer to as "minimum dissimilarity pruning".

2.2.1 Minimum dissimilarity pruning

In this stage, we define a measure of the dissimilarity between sibling nodes, and then collapse these nodes if this measure is lower than a certain threshold. We first consider the maximal tree T_0 obtained in the previous stage. Let t_l and t_r be a pair of terminal nodes that share the same direct ascendant. We then define (in population terms) the random variables $W_{l(r)} = D(X_{t_l}, sop(X_{t_r}))$, (where sop(Z) stands for the support of the random variable Z) as the Euclidian distance between the random elements of X_{t_l} and the support of X_{t_r} and respectively $W_{r(l)} = D(sop(X_{t_l}), X_{t_r})$. For each of them we define a dissimilarity measure

$$\Delta_{l(r)} = \int_0^\delta q_\nu(W_{l(r)}) d\nu,$$

$$\Delta_{r(l)} = \int_0^\delta q_\nu(W_{r(l)}) d\nu,$$

(4)

where $q_{\nu}(W_{lr})$ stands for the quantile function, $(\mathbb{P}(W_{l(r)} \leq q_{\nu}) = \nu)$ and δ is a proportion, $\delta \in (0, 1)$.

Finally we consider as a measure of dissimilarity between the sets t_l and t_r

$$\Delta_{lr} = max\{\Delta_{l(r)}, \Delta_{r(l)}\}.$$

If $\Delta_{lr} < \epsilon$, we prune at the node t, i.e. we replace t_l and t_r by $t_l \cup t_r$ in the partition.

Observe that since $\Delta_{l(r)}$ and $\Delta_{r(l)}$ are averages of small quantiles of $D(X_{t_l}, sop(X_{t_r}))$ and of $D(X_{t_r}, sop(X_{t_l}))$ respectively, Δ_{lr} can be thought as "a more resistant version" of the distance between the supports of the random vectors X_{t_l} and X_{t_r} .

We consider the natural plug-in estimate for the dissimilarity measure Δ , that is defined as follows. Let n_l (resp. n_r) be the size of \hat{t}_l (resp. \hat{t}_r). We may then consider for every $x_i \in \hat{t}_l$ and $y_j \in \hat{t}_r$, the sequences, $\tilde{d}_i = \min_{y \in \hat{t}_l} d(x_i, y)$, $\tilde{d}_j = \min_{x \in \hat{t}_r} d(x, y_j)$ and their ordered versions, denoted as d_i and d_j . For $\delta \in [0, 1]$, let

$$\bar{d}_l^{\delta} = \frac{1}{\delta n_l} \sum_{i=1}^{\delta n_l} d_i, \quad \bar{d}_r^{\delta} = \frac{1}{\delta n_r} \sum_{j=1}^{\delta n_r} d_j$$

We compute the dissimilarity between t_l and t_r to be,

$$d^{\delta}(l,r) = d^{\delta}(t_l,t_r) = max(\bar{d}_l^{\delta},\bar{d}_r^{\delta}),$$

and at each stage of the algorithm the leaves, t_l and t_r , are merged into the ascendant node t if $d^{\delta}(l,r) \leq \epsilon$ where $\epsilon > 0$. The dissimilarity pruning makes use of the two parameters δ and ϵ , which we hereafter refer to as "mindist".

2.2.2 Joining

The aim of the joining step is to aggregate those nodes that do not share the same direct ascendent. The criterion used in the joining step is the same as the used on the pruning one without the restriction of being sibling nodes. The need of this step is illustrated in Figure 1.



Figure 1: The lower group is split on the first stage and it cannot be merged in the pruning stage, but will be when using the joining step.

Here, all pairs of terminal nodes t_i and t_j are compared by computing, $d^{\delta}(i, j)$. As in standard hierarchical clustering, pairs of terminal nodes are successively aggregated starting from the pair i, j with minimum $d^{\delta}(i, j)$ value, thereby producing one less cluster at each step. We consider two different stopping rules for the joining procedure, which correspond to the two cases where the number of clusters k is either known or unknown. We denote by m the number of leaves of the pruned tree. If k is known, the following step is repeated until $m \leq k$:

• For each pair of values (i, j), $1 \leq i < j \leq m$, let $(\tilde{i}, \tilde{j}) = \arg \min_{i,j} \{ d^{\delta}(i, j) \}$. Replace $t_{\tilde{i}}$ and $t_{\tilde{j}}$ by its union $t_{\tilde{i}} \cup t_{\tilde{j}}$, put m = m - 1 and proceed.

If k is unknown,

• if $d_{\tilde{i}\tilde{j}}^{\delta} < \eta$ replace $t_{\tilde{i}}$ and $t_{\tilde{j}}$ by its union $t_{\tilde{i}} \cup t_{\tilde{j}}$, where $\eta > 0$ is a given constant, and continue until this condition does not apply.

In the first case, the stopping criterion is the number of clusters, while in the second case a threshold value of η for $d^{\delta}(i, j)$ must be specified.

2.3 CUBT and k-means

In the following section we discuss informally the circumstances in which both our procedure and the well-known k-means algorithm should produce a reasonably good set of results. We shall consider those cases where there are "nice groups" that are strictly separated. More precisely, let A_1, \ldots, A_k be disjoint connected compact sets on \mathbb{R}^p such that $A_i = \overline{A_i^0}$ for $i = 1, \ldots, k$, and $\{P_i : i = 1, \ldots, k\}$ their probability measures on \mathbb{R}^p with supports $\{A_i : i = 1, \ldots, k\}$.

A typical case may be obtained by defining a random vector \mathbf{X}^* with a density f and then considering the random vector $\mathbf{X} = \mathbf{X}^* | \{f > \zeta\}$ for a positive level set ζ , as in a number of hierarchical clustering procedures.

An admissible family for CUBT is the family of sets A_1, \ldots, A_k such that there exist another family of disjoint sets B_1, \ldots, B_k that are built up as the intersection of a finite number of half-spaces delimited by hyperplanes that are orthogonal to the coordinate axis that satisfying $A_i \subset B_i$. In contrast, the k-means algorithm is defined through the vector of centres $(c_1 \ldots, c_k)$ that minimise $\mathbb{E}(\min_{j=1,\ldots,k} \|\mathbf{X} - c_j\|)$. Associated with each centre c_j is the convex polyhedron S_j of all points in \mathbb{R}^p that are closer to c_j than to any other center, called the Voronoi cell of c_j . The sets in the partition S_1, \ldots, S_k are the population clusters for the *k*-means algorithm. The population clusters for the *k*-means algorithm are therefore defined by exactly *k* hyperplanes in an arbitrary position.

Then, an admissible family for the k-means algorithm will be a family of sets A_1, \ldots, A_k that can be separated by exactly k hyperplanes.

Although the hyperplanes can be in an arbitrary position, no more than k of them can be used.

It is clear that in this sense CUBT is much more flexible than the k-means algorithm, because the family of admissible sets is more general. For example, k-means will necessarily fail to identify nested groups, while CUBT will not.

Another important difference between k-means and CUBT is that our proposal is less sensitive to small changes in the parameters that define the partition. Effectively, small changes in these will produce small changes in the partition. However, small changes in the centres $(c_1 \ldots, c_k)$ that define the k-means partition can produce significant changes in the associated partition as given by the Voronoi cells.

Figure 2 show an example where CUBT has a good performance and clearly k-means does not work.



Figure 2: (a) CUBT cluster allocation. (b) k-means cluster allocation.

3 Consistency of CUBT

In the following section we present some theoretical results on the consistency of our algorithm. We first prove an important property, that of the monotonicity of the deviance with increasing tree size.

A simple equivalent characterisation of the function R(t) is given in the following Lemma.

Lemma 3.1. Let t_l and t_r be disjoint compact sets on \mathbb{R}^p and $\mu_s = \mathbb{E}(X_{t_s})$, s = l, r. If $t = t_l \cup t_r$ we have,

$$R(t) = R(t_l) + R(t_r) + \frac{\alpha_{t_l} \alpha_{t_r}}{\alpha_t} \|\mu_l - \mu_r\|^2.$$
 (5)

The proof is given in the Appendix.

Remark 3.1. Monotonicity of the function R(.) and geometric interpretation. Observe that Lemma (3.1) entails that for all disjoint compact sets t_l, t_r and $t = t_l \cup t_r$, the function R(.) is monotonic in the sense that, $R(t) \ge R(t_l) + R(t_r)$. Moreover, R(t)will be close to $R(t_l) + R(t_r)$ when the last term on the right hand side of equation (5) is small. This will occur either if one of the sets t_l and t_r has a very small fraction of the mass of t and/or if the centers of the two subsets t_l, t_r are very close together. In either cases we do not want to split the set t.

The following results show the consistency of the empirical algorithm in relation to its population version. We begin with the splitting algorithm and then follow this by pruning and joining.

Theorem 3.1. Assume that the random vector \mathbf{X} has distribution P and a density f that fulfils the condition that $||x||^2 f(x)$ is bounded. Let $\mathbf{X}_1, \ldots, \mathbf{X}_n$ be iid random vectors with the same distribution as \mathbf{X} and denote by P_n the empirical distribution of the sample $\mathbf{X}_1, \ldots, \mathbf{X}_n$.

Let $\{t_{1n}, \ldots, t_{m_nn}\}$ be the empirical binary partition obtained from the forward empirical algorithm, and $\{t_1, \ldots, t_m\}$ the population version. We then find that ultimately, i.e. there exists n_0 such that for $n > n_0$, $m_n = m$ and each pair $(i_{jn}, a_{jn}) \in \{1, \ldots, p\} \times \mathbb{R}$ in the determination of the empirical partition converges a.s. to the corresponding one $(i_j, a_j) \in \{1, \ldots, p\} \times \mathbb{R}$ for the population values. In particular, it implies that, $\lim_{n\to\infty} \sum_{i=1}^m P(t_{in}\Delta t_i) = 0$, where Δ stands for the symmetric difference.

The proof is given in the Appendix.

Theorem 3.2. Let $\{t_{1n}^*, \ldots, t_{knn}^*\}$ be the final empirical binary partition obtained after the forward and backward empirical algorithm has been applied, and $\{t_1^*, \ldots, t_k^*\}$ be the population version. Under the assumptions of Theorem 3.1 we ultimately have that $k_n = k$ $(k_n = k \text{ for all } n \text{ if } k \text{ is known})$, and $\lim_{n\to\infty} \sum_{i=1}^k P(t_{in}^*\Delta t_i^*) = 0$.

The proof is given in the Appendix.

4 Some experiments.

In the following Section, we present the results of a simulation in which we tested our method using four different models. We consider separately the cases where the number of groups is known and where it is unknown. If the number of groups is known we compare the results obtained with those found using the k-means algorithm. It is well known the performance of the k-means algorithm depends strongly on the position of the centroids initially used to start the algorithm, and a number of different methods have been proposed to take account of this effect (see Steinley, [23]). We herein follow the recommendations made in this last reference, and consider ten random initialisations, keeping the one with the minimum within-cluster sum of squares given by, $\sum_{i=1}^{n} \sum_{j=1}^{k} ||\mathbf{X}_i - c_j||^2 \mathcal{I}_{\{\mathbf{X}_i \in G_j\}}$, where G_j is the j-th group and c_j is the corresponding center. We refer to this version as k-means(10).

We also compare our results with the well known model-based clustering procedure proposed by Fraley and Raftery ([7], [8]). The MCLUST algorithm assumes that the sample has been generated from a mixture of G normal distributions, with ellipsoidal covariance matrices with variable shape, volume and orientation, and estimates the parameters of each population considering an EM algorithm and the most appropriate model is selected by means of the Bayesian Information Criteria. This procedure is versatile and it can be applied whether the number of groups is stated or not, hence we report the results for both cases.

We consider the algorithm DBSCAN (Density Based Spatial Clustering of Applications with Noise, [6]) which is suitable for discovering arbitrarily shaped clusters, since clusters are defined as dense regions separated by low-density regions. This algorithm estimates the number of clusters, as most of the density based algorithms do, hence we compare the results with those where the number of groups is unknown. The algorithm depends on two parameters, the number of objects in a neighborhood of an object, our input was 5, and the neighborhood radius, the authors, Ester et al [6], recommend to avoid putting an input if this parameter is unknown.

4.1 Simulated data sets

We herein consider four different models (M1 - M4) having different number of groups and dimensions. For each model clusters have equal number of observations.

- M1. Four groups in dimension 2 with 100 observations per group. The data are generated using : $N(\mu_i, \Sigma)$, i = 1, ..., 4, and distributed with centers (-1, 0), (1, 0), (0, -1), (0, 1) and covariance matrix $\Sigma = \sigma^2 I d$ with $\sigma = 0.11, 0.13, 0.15, 0.17, 0.19$. Figure 3 (a) gives an example of the data generated using model M1 with $\sigma = 0.19$.
- M2. Ten groups in dimension 5 with 30 observations per group. The data are generated using $N(\mu_i, \Sigma), i = 1, ..., 10$. The means μ_i of the first five groups, are the vectors of the canonical basis e_1, \ldots, e_5 respectively, while the centers of the five remaining groups are $\mu_{5+i} = -e_i$, $i = 1, \ldots, 5$. In all cases, the covariance matrix is $\Sigma = \sigma^2 Id$, $\sigma = 0.11, 0.13, 0.15, 0.17$. Figure 3 (b) gives an example of data generated using model M4 with $\sigma = 0.19$ projected on the two first dimensions.
- M3. Two groups in dimension 2 with 150 observations per group. The data are uniformly generated on two concentric rings. The radius of the inner ring is between 50 and 80 and the radius of the outer ring is between 200 and 230. Figure 3 (c) gives an example of data generated using model M3.
- M4. Three groups in dimension 50 with 25 observations per group. The data are generated using $N(\mu_i, \Sigma)$, i = 1, 2, 3. The mean of the first group is $(0.1, \ldots, 0.1)$, the second group is centered at the origin and the last group has mean $(-0.1, \ldots, -0.1)$. In all cases, the covariance matrix is $\Sigma = \sigma^2 Id$, $\sigma = 0.03$ or 0.05.



Figure 3: (a) Scatter plot corresponding to M1 for $\sigma = 0.19$ (b) Two-dimensional projection scatter plot corresponding to M2 for $\sigma = 0.19$ (c) Scatter plots corresponding to M3

4.2 Tuning the method

We performed M = 100 replicates for each model. When k is given, we compare the results with those obtained using k-means, k-means (10) and MCLUST. Otherwise, if k is unknown we compared the results with DBSCAN and MCLUST. In order to apply CUBT we must fix the values for the parameters involved at each stage in the algorithm: for the maximal tree we used minsize = 5, 10 or 15 and mindev = 0.7 or 0.9; for the pruning stage mindist = 0.3 or 0.5 and $\delta = 0.2, 0.4$ or 0.6 for the joining stage. For the cases where the number of clusters is stated, it is important to note that even though in almost every case the number of terminal nodes of the maximal tree is bigger than the number of groups, there is no warranty that this will happen. Then, if we are in that case we reduce *mindev* in decreasing order, $0.6, 0.5, \ldots$ until the maximal tree has at least k terminal nodes. For the cases where the number of clusters is not stated we consider the same parameters as in the previous case. To determine the number of clusters we must choose a threshold η . We consider the distances $d^{\delta}_{i,\tilde{j}}$ defined in Section 2.2.2 that correspond to the tree that is the output of the pruning step, and fix η as a low quantile of $d_{\tilde{i},\tilde{j}}^{\delta}$. Heuristically, low quantiles of $d_{\tilde{i},\tilde{j}}^{\delta}$ correspond to terminal nodes whose observations belong to the same clusters. The quantiles of $d^{\delta}_{i,\tilde{j}}$ that determine η chosen for M1 to M4 were 0.2, 0.08, 0.25 and 0.15 respectively.

Because we use synthetic data sets, we know the actual label of each observation, and it is thus reasonable to measure the goodness of a partition by computing "the number of misclassified observations", which is analogous to the misclassification error for supervised classification procedures. We denote the original clusters $r = 1, \ldots, R$. Let $y_1, \ldots y_n$ be the group label of each observation, and $\hat{y}_1, \ldots \hat{y}_n$ the class label assigned by the clustering algorithm. Let Σ be the set of permutations over $\{1, \ldots, R\}$. The misclassification error may then be expressed as:

$$MCE = \min_{\sigma \in \Sigma} \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}_{\{y_i \neq \sigma(\hat{y}_i)\}}.$$
(6)

If the number of clusters is large, the assignment problem may be computed in polynomial time using Bipartite Matching and the Hungarian Method, [17]. We use this algorithm only for M2 that has ten groups.

4.3 Results

First, we analyze the results of the simulation when the number of clusters is known. Table 1 shows the results obtained in our simulations. Except for Model 3, we varied the values of σ to reflect the degree of overlapping between classes. In order to study the robustness of the method with respects to the parameters selection we considered 36 different parameters configurations for our procedure. Since the results were practically the same, we present for each model the results for the best and the worst clustering allocation. We report the misclassification error obtained for each clustering procedure, namely CUBT for the best performance (CUBT (B)) and for the worst performance (CUBT (W)), k-Means, k-Means(10) and MCLUST (if k is given). As expected for the first two models k-means(10) and MCLUST have a good performance, but both of them fail for M3. For the last model MCLUST performs poorly in both cases, k-means has an acceptable performance in both cases and k-means (10) always achieves a perfect cluster allocation. The results obtained by CUBT using different values for the parameters are practically the same, and in almost all cases the results for CUBT lie between those of k-Means and k-Means(10), except for Model 3 where it has a better performance and for M4 when $\sigma = 0.05$ (where the misclassification rate of CUBT(W) is larger than the misclassification rate of k-means). If we compare the results with those of MCLUST we may say that for M1 they have a similar performance and that for M2, M3 and M4 the performance of CUBT is better.

Sigma (σ)	CUBT (W)	CUBT (B)	k-Means	k-Means(10)	MCLUST	
	Model 1					
0.11	0	0	0.12	0	0	
0.13	7e - 03	0	0.12	0	0	
0.15	1e - 03	1e - 04	0.11	0	0	
0.17	1e - 03	2e - 04	0.07	5e - 05	3e - 05	
0.19	2e - 03	3e - 04	0.06	3e - 04	2e - 04	
Model 2						
0.7	0	0	0.11	0.01	0.04	
0.75	0	0	0.10	0.01	0.05	
0.8	4e - 04	2e - 04	0.10	0.01	0.06	
0.85	0.004	0.002	0.08	0.01	0.07	
0.9	0.05	0.04	0.07	0.01	0.08	
Model 3						
	0	3e - 04	0.47	0.47	0.25	
Model 4						
0.03	0	0	0.11	0	0.65	
0.05	0.16	0.05	0.12	0	0.65	

Table 1: Simulation results for models M1 to M4.

Now we proceed to analyze the results of the simulation when k is unknown. In Table 2 we report the number of times that the procedure chooses the correct number of groups. The number of misclassified observations for CUBT and MCLUST, when the number of

clusters is chosen properly, is the same as in the previous analysis. CUBT has a very good performance for M1, M3 and the first case of M4, and for the others situations the best performance is very good but it strongly depends on the choices of the other parameters. However, is important to note that in those cases identifying correctly the clusters is a very difficult problem and the other methods were not able to do it. For M2 if $\sigma = 0.7$ DBSCAN in nine opportunities identifies ten groups (in those cases all the observations are classified correctly) and in the rest of the replicates it identifies fewer groups, if the group overlapping is bigger it is not able to separate the groups, for $\sigma = 0.75$ it always finds less than five groups and on the rest of the cases it always finds only one group. For M1 DBSCAN has a very good performance identifying the number of groups and in those cases there are not misclassified observations except for the case of $\sigma = 0.19$ that the misclassification rate is 0.027. For M3 and M4 whenever it identifies correctly the number of groups it also allocates the observations in the right group. MCLUST has an outstanding performance for M1, but it fails in all the other models. For M2 it always identifies one group and for M3 97 times finds nine groups and the rest of the times eight clusters. Finally for M4 even though it always finds three groups it fails in the allocation of the observations, which is consistent with the results found when the number of clusters was previously stated.

Sigma (σ)	CUBT (W)	CUBT (B)	DBSCAN	MCLUST	
Model 1					
0.11	81	85	92	99	
0.13	75	84	81	100	
0.15	76	87	74	100	
0.17	79	92	43	100	
0.19	85	88	38	100	
Model 2					
0.7	24	90	9	0	
0.75	34	90	0	0	
0.8	42	90	0	0	
0.85	58	83	0	0	
0.9	51	78	0	0	
Model 3					
	94	98	76	0	
Model 4					
0.03	100	100	100	100	
0.05	2	98	26	100	

Table 2: Number of times that the different procedures choose the correct number of groups for M1 to M4.

4.4 A comparison between CART and CUBT

In the following simple example, we compare the tree structure obtained by CART using the complete training sample (observations plus the group labels) with that obtained by CUBT considering only the training sample without the labels. We generated three sub-populations in a two-dimensional variable space. The underlying distribution of the vector $\mathbf{X} = (X_1, X_2)$ is a bivariate normal distribution in which the variables X_1 and X_2 are independent, with their distributions for the three groups being given by $X_1 \sim$ $\mathcal{N}(0, 0.03)$, $\mathcal{N}(2, 0.03)$, $\mathcal{N}(1, 0.25)$, $X_2 \sim \mathcal{N}(0, 0.25)$, $\mathcal{N}(1, 0.25)$, $\mathcal{N}(2.5, 0.03)$.

The data were then rotated through $\pi/4$. One of the difficulties was that the optimal partitions are not parallel to the axis. Figure 4 shows the partitions obtained using CART and CUBT, for a single sample of size 300. We performed 100 replicates, and in each case generated a training sample of size 300, whereby every group was of the same size. We then computed the "mean misclassification rate" with respect to the true labels. For CUBT, the value was 0.09, while for CART there were no classification errors because we used the same sample both for growing the tree and for classifying the observations.



Figure 4: Plot of the partitions of the space for a generated data set. The solid lines indicate the partition for CUBT and the dashed lines the partition for CART

In order to compare our procedure with the traditional CART classification method, we obtained the binary trees for CART and CUBT. Both trees are shown in Figure 5. It is noteworthy that the two structures are exactly the same, and that the different cutoff values in the different branches may be understood with the aid of Figure 4, which corresponds to the same data set.

5 A real data example. European Jobs

In the following example, the data set describes the percentage of workers employed in different sectors of economic activity for a number of European countries in the year 1979



Figure 5: <u>Left:</u> Tree corresponding to CUBT . <u>Right:</u> Tree corresponding to CART. In both cases the left-hand branches indicate the smaller value of the partitioning variable.

(this data set can be obtained on the website http://lib.stat.cmu.edu/DASL/Datafiles

/EuropeanJobs.html). The categories are: agriculture (A), mining (M), manufacturing (MA), power supplies industries (P), construction (C), service industries (SI), finance (F), social and personal services (S) and transportation and communication (T). It is important to note that these data were collected during the cold war. The aim was to allocate the observations to different clusters, but the number of clusters is unknown. We must therefore study the data structure for a range of different numbers of clusters. We first consider a four-group structure. In this case, a single variable (the percentage employed in agriculture) determines the tree structure. The four groups are shown in Table 3, and the corresponding tree is plotted on the top panel of Figure 6. In the tree, the highest value of A corresponds to Turkey, which is an outlier and conforms to a single cluster of observations, which may be explained in terms of its social and territorial proximity to Africa. The countries that make up groups 2 and 3 are those that were either under communist rule or those that were experiencing varying degrees of political upheaval; Spain, for example was making adjustments after the end of Franco's regime. The countries of Group 2 were poorer than those of Group 3. Finally, Group 4 had the lowest percentage of employment in agriculture, and the countries in this group were the most developed and were not under communist rule, with the exception of East Germany. Using k-means we get the following clusters. Turkey and Ireland are isolated in one group each, Greece, Portugal, Poland, Romania and Yugoslavia form another group, and the remaining countries form a fourth cluster.

If we use five clusters instead of four, the main difference is that Group 4 of the original partition is divided into two subgroups (4 and 5), and the variables that explain these partitions are the percentages employed in mining and agriculture. The other groups of the original partition remain stable.

If a five-cluster structure is considered via the k-means algorithm, Turkey and Ireland are then each isolated in single groups, and Greece, Portugal, Poland, Romania and Yugoslavia form another group, as in the four-cluster structure. Switzerland and East and West Germany make up a new cluster.

Group 1	Group 2	Group 3	Group 4
Turkey	Greece	Ireland	Belgium
	Poland	Portugal	Denmark
	Romania	Spain	France
	Yugoslavia	Bulgaria	W. Germany
		Czechoslovakia	E. Germany
		Hungary	Italy
		USSR	Luxembourg
			Netherlands
			United Kingdom
			Austria
			Finland
			Norway
			Sweden
			Switzerland

Table 3: CUBT clustering structure using four groups.





6 Concluding Remarks

We have herein presented a new clustering method called CUBT which shares some ideas with the well known classification and regression trees, defining clusters in terms of binary rules over the original variables. Like CART, our method may be very attractive and useful in a number of practical applications. Because the tree structure makes use of the original variables, it helps to determine those variables that are important in the conformation of the clusters. Moreover, the tree allows the classification of new observations. In our

Group 1	Group 2	Group 3	Group 4	Group 5
Turkey	Greece	Ireland	Belgium	W. Germany
	Poland	Portugal	Denmark	E. Germany
	Romania	Spain	France	Switzerland
	Yugoslavia	Bulgaria	Italy	
		Czechoslovakia	Luxembourg	
		Hungary	Netherlands	
		USSR	United Kingdom	
			Austria	
			Finland	
			Norway	
			Sweden	

Table 4: CUBT clustering with five groups.

approach, a binary tree is obtained in three stages. In the first stage, the sample is split into two sub-samples, thereby reducing the heterogeneity of the data within the new sub-samples according to the objective function $R(\cdot)$. The procedure is then applied recursively to each sub-sample. In the second and third stages, the maximal tree obtained at the first stage is pruned using a dissimilarity criterion, first applied to the adjacent nodes and then to all the terminal nodes. The algorithm is simple and requires a reasonable computation time. There are no restrictions on the dimension of the data used. Our method is consistent under a set of quite general assumptions, and produces quite good results with the simulated examples that we considered here, as well as for an example that used real data.

The algorithm depends on several parameters, and an optimal way to choose them is beyond the scope of this paper. We herein propose some advice in order choose them in practice. The splitting stage depends on two parameters, *mindev* and *minsize*. The later, mindev $\in (0,1)$ represents the percentage of the deviance of the whole data set $(R(\mathcal{S}))$ the algorithm requires to split a group (if the reduction of its deviance is less than $mindev \times R(\mathcal{S})$ the group is split in two subgroups). Our experience indicates that values between 0.7 and 0.9 give a sufficiently large partition. The former indicates the minimum cluster size that the user admits, if he has some information beforehand it could be provided and taken into account by the algorithm otherwise the default value should be 1. The pruning step includes also two parameters δ and ϵ . In population terms, it suffices that ϵ be smaller than the distance among the supports of two disjoint clusters. If the user could not provide an input for this parameter the default could be zero, which correspond to skip the pruning step and go directly to the joining step, in that case one would probably obtain a larger tree, but the final clustering allocation would not change and the results given in Theorem 3.2 still hold. The parameter δ is just a way to deal with some possible presence of outliers and as a default $\delta = 0.2$ can be used. Finally, if the number of clusters is given the final stage does not require any parameter, otherwise one parameter, η , should be provided. The way we suggest to choose this parameter is given in Section 4.2.

A more robust version of our method could be developed by substituting the objective function $cov(X_T)$ in equation (1) for a more robust covariance functional $robcov(X_T)$ (see

e.g. Maronna et al. [15] Chapter 6 for a review), and then proceeding in the same manner as described herein. However, a detailed investigation of this alternative approach lies beyond the scope of the present study.

Even though we focussed our paper on the treatment of interval data, the procedure can be extended to other settings as long as a covariance structure and a dissimilarity measure can be defined. For example, one could use Gini's definition of variance for categorical data, ([11]) and the Mean Character difference as a measure of dissimilarity. Other measures of dissimilarity for categorical data can be found in Gan et al.[9]. However, a deeper study of these settings are beyond the scope of this paper.

Acknowledgments 7

Ricardo Fraiman and Marcela Svarc have been partially supported by Grant pict 2008-0921 from ANPCyT (Agencia Nacional de Promocin Científica y Tecnológica), Argentina.

Appendix 8

Proof of Lemma 3.1 8.1

We first observe that because t_l and t_r are disjoint, $\mathbb{E}(X_{t_l \cup t_r}) = \gamma \mu_l + (1 - \gamma)\mu_r$, where $\gamma = P(\mathbf{X} \in t_l | \mathbf{X} \in t_l \cup t_r)$. Given that $j = 1, \dots, p$, we use $M_{2i}^{(j)} = \int_{t_i} x(j)^2 dF(x)$, i = l, r, where F stands for the distribution function of the vector \mathbf{X} . It then follows that $\mathbb{E}(X_{t_l \cup t_r}(j)^2) = \gamma M_{2l}^{(j)} + (1 - \gamma) M_{2r}^{(j)}$, and therefore that

$$var(X_{t_l \cup t_r}(j)) = \gamma var(X_{t_r}(j)) + (1 - \gamma)var(X_{t_r}(j)) + \gamma(1 - \gamma)(\mu_l(j) - \mu_r(j))^2.$$

By summing the terms in j we get the desired result.

8.2 Proof of Theorem 3.1

Let \mathbb{T} be the family of polygons in \mathbb{R}^p with faces orthogonal to the axes, and fix $i \in$ $\{1,\ldots,p\}$ and $t \in \mathbb{T}$. For $a \in \mathbb{R}$ denote by $t_l = \{x \in t : x(i) \leq a\}$ and $t_r = t \setminus t_l$. We define $r(t, i, a) = R(t) - R(t_l) - R(t_r)$ and $r_n(t, i, a) = R_n(t) - R_n(t_l) - R_n(t_r)$, the corresponding empirical version.

We start showing the uniform convergence

$$\sup_{a \in \mathbb{R}} \sup_{t \in \mathbb{T}} |r_n(t, i, a) - r(t, i, a)| \to 0 \quad a.s.$$
(7)

By Lemma 3.1,

$$\alpha_t r(t, i, a) = \alpha_{t_l} \alpha_{t_r} \|\mu_l(a) - \mu_r(a)\|^2,$$
(8)

where $\alpha_A = P(\mathbf{X} \in A)$ and $\mu_j(a) = \mathbb{E}(X_{t_j}), j = l, r$. Then, the pairs (i_{jn}, a_{jn}) and (i_j, a_j) are the arguments that maximise the right-hand side of equation (8) with respect to the measures P_n and P respectively. We observe that the right-hand side of equation (8) equals

$$\alpha_{t_r} \int_{t_l} \|x\|^2 dP(x) + \alpha_{t_l} \int_{t_r} \|x\|^2 dP(x) - 2\langle \int_{t_l} x dP(x), \int_{t_r} x dP(x) \rangle.$$
(9)

It order to prove equation (7) it is sufficient to show that:

- 1. $\sup_{a \in \mathbb{R}} \sup_{t \in \mathbb{T}} |P_n(t_j) P(t_j)| \to 0 \quad a.s. \quad j = l, r$
- 2. $\sup_{a \in \mathbb{R}} \sup_{t \in \mathbb{T}} |\int_{t_j} ||x||^2 dP_n(x) \int_{t_j} ||x||^2 dP(x)| \to 0$ a.s. j = l, r
- 3. $sup_{a\in\mathbb{R}}\sup_{t\in\mathbb{T}} |\int_{t_j} x(i)dP_n(x) \int_{t_j} x(i)dP(x)| \to 0 \ a.s. \ j = l, r, i = 1, \dots, p.$

Since \mathbb{T} is a Vapnik–Chervonenkis class, we have that (i) holds. Now observe that the conditions for uniform convergence over families of sets still hold if we are dealing with signed finite measures. Therefore if we consider the finite measure $||x||^2 dP(x)$ and the finite signed measure given by x(i)dP(x) we also have that (ii) and (iii) both hold.

Since

$$\lim_{a \to \infty} \alpha_{t_l} \alpha_{t_r} \|\mu_l(a) - \mu_r(a)\|^2 = \lim_{a \to -\infty} \alpha_{t_l} \alpha_{t_r} \|\mu_l(a) - \mu_r(a)\|^2 = 0,$$

we have that $\inf \{ \operatorname{argmax}_{a \in \mathbb{R}} r_n(t, i, a) \} \to \inf \{ \operatorname{argmax}_{a \in \mathbb{R}} r(t, i, a) \}$ a.s.

In the first step of the algorithm, $t = \mathbb{R}^p$ and we obtain $i_{n1} = i_1$ for n large enough and $a_{n1} \to a_1$ a.s. In the next step, we observe that the empirical procedure begins to work with t_{nl} and t_{nr} , while the population algorithm will do so with t_l and t_r . However, we have that

$$\sup_{a \in \mathbb{R}} |r_n(t_{nj}, i, a) - r(t_j, i, a)| \le \sup_{a \in \mathbb{R}} \sup_{t \in \mathbb{T}} |r_n(t_{nj}, i, a) - r(t_{nj}, i, a)| + \sup_{a \in \mathbb{R}} |r(t_{nj}, i, a) - r(t_j, i, a)|,$$
(10)

for j = l, r.

We already know that the first term on the right hand side of equation(10) converges to zero almost surely. In order to show that the second term also converges to zero, it is sufficient to show that

1.
$$\sup_{a \in \mathbb{R}} |P(t_{nj}) - P(t_j)| \to 0$$
 a.s. $j = l, r$
2. $\sup_{a \in \mathbb{R}} |\int_{t_j} ||x||^2 dP(x) - \int_{t_{nj}} ||x||^2 dP(x)| \to 0$ a.s. $j = l, r$
3. $\sup_{a \in \mathbb{R}} |\int_{t_j} x(i) dP(x) - \int_{t_{nj}} x(i) dP(x)| \to 0$ a.s. $j = l, r, i = 1, ..., p,$

which follows from the assumption that $||x||^2 f(x)$ is bounded. This concludes the proof since $minsize/n \to \tau$.

8.3 Proof of Theorem 3.2

We need to show that we have consistency in both steps of the backward algorithm.

(i) Convergence of the pruning step. Let $\{t_{1n}^*, \ldots, t_{mn}^*\}$ be the output of the forward algorithm. The pruning step partition of the algorithm converges to the corresponding population version from

- the conclusions of Theorem 3.1.
- the fact that the random variables W_{lr} \tilde{d}_l, \tilde{d}_r are positive.
- the uniform convergence of the empirical quantile function to its population version.
- the Lebesgue dominated convergence Theorem.

The proof of (ii) is mainly the same as that for (i).

References

- Andreopoulos, B., An, A. and Wang, X. (2007). "Hierarchical Density-Based Clustering of Categorical Data and a Simplification". Advances in Knowledge Discovery and Data Mining, 11–22.
- [2] Basak, J. and Krishnapuram, R. (2005). "Interpretable Hierarchical Clustering by Construction an Unsupervised Decision Tree". *IEEE Transactions on Knowledge and Data Engineering*, **17** 121-132.
- [3] Blockeel, H., De Raedt, L. and Ramon, J. (1998). "Top-down induction of clustering trees, Proceedings of the 15th International Conference on Machine Learning". ed. Morgan Kaufmann, 55–63.
- [4] Breiman, L., Friedman, J., Stone, C. J. and Olshen, R. A. (1984), "Classification and Regression Trees", ed. Chapman & Hall/CRC.
- [5] Chavent, M., Guinot, C., Lechevallier, Y. and Tenenhaus, M. (1999). "Méthodes Divisives de Classification et Segmentation Non Supervisée: Recherche d'une Typologie de la Peau Humanine Saine". *Rev. Statistique Appliquée, XLVII* 87–99.
- [6] Ester, M., Kriegel, H.-P., Sander, J. and Xu, X.(1996). "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise". Proc. 2nd Int. Conf. on Knowledge Discovery and Data Mining. Portland, OR, 226-231.
- [7] Fraley, C. and Raftery, A.E. (2002). "Model-based clustering, discriminant analysis, and density estimation". *Journal of the American Statistical Association*, **97**, 611–631.
- [8] Fraley, C. and Raftery, A. E. (2009). "MCLUST Version 3 for R: Normal Mixture Modeling and Model-based Clustering", Technical Report No. 504, Department of Statistics, University of Washington.

- [9] Gan, G., Ma, C. and Wu, J. (2007). "Clustering Data: Theory, Algorithms and Applications", Springer.
- [10] García-Escudero, L. Gordaliza, A. Matrán, C. and Mayo-Iscar, A. (2010). "A review of robust clustering methods". Advances in Data Analysis and Classification, 4, 2, 89-109.
- [11] Gini, C.W.(1912). "Variability and Mutability, contribution to the study of statistical distributions and relations". Studi Economico-Giuridici della R. Universita de Cagliari.
- [12] Gnanadesikan, R., Kettering, J. R. and Maloor, S. (2007). "Better Alternatives to Current Methods of Scaling and Weighting Data for Cluster Analysis". *Journal of Statistical Planning and Inference* 137, 3483-3496.
- [13] Hastie, T., Tibshirani, R. and Friedman, J. H. (2003). "The Elements of Statistical Learning", Springer.
- [14] Liu, B., Xia, Y. and Yu, P. S. (2000). "Clustering through decision tree construction", CIKM '00: Proceedings of the ninth international conference on Information and knowledge management, ed. ACM, New York, NY, USA, 20–29.
- [15] Maronna, R. A., Martin, D. R. and Yohai, V. J. (2006). "Robust Statistics: Theory and Methods", Wiley Series in Probability and Statistics, Wiley.
- [16] Oh, M. S. and Raftery, A. (2007., "Model-Based Clustering With Dissimilarities: A Bayesian Approach", Journal of Computational and Graphical Statistics 16(3), 559-585.
- [17] Papadimitriou, C. and Steiglitz, K. (1982). "Combinatorial Optimization: Algorithms and Complexity". Englewood Cliffs: Prentice Hall.
- [18] Peña, D. and Prieto F. J. (2001). "Cluster Identification using Projections". Journal of American Statistical Association, 96, 456, 1433-1445.
- [19] Peña, D. and Prieto F. J. (2006). "A Projection Method for Robust Estimation and Clustering in Large Data Set" in Data Analysis, Classification and the Forward Search, Zani S. Cerioli A. Riani M. Vichi M. eds, Springer Verlag, Berlin.
- [20] Questier F., R. Put R., Coomans D., Walczak B. and Vander Heyden Y. (2005). "The use of CART and multivariate regression trees for supervised and unsupervised feature selection". *Chemometrics and Intelligent Laboratory Systems*, **76** 1 45–54.
- [21] Smyth C., Coomans D., Everingham Y. and Hancock T. (2006). "Auto-associative Multivariate Regression Trees for Cluster Analysis", *Chemometrics and Intelligent Lab*oratory Systems, 80 1 120–129.
- [22] Smyth C., Coomans D. and Everingham Y. (2006). "Clustering noisy data in a reduced dimension space via multivariate regression trees", *Pattern Recognition*, **39** 3 424–431.

- [23] Steinley D. and Brusco M. (2007). "Initializing K-means Batch Clustering: A Critical evaluation of Several Techniques". *Journal of Classification*, **24** 99–121.
- [24] Tryon, R.C. (1939). "Cluster Analysis", McGraw-Hill.
- [25] Walther G. (2010). "Optimal and fast detection of spatial clusters with scan statistics". The Annals of Statistics, 24 N 2 1010–1033.